⍬

AD-A219 636

Report No. 7176

Detecting Black Holes in Packet-Radio Networks (SRNTN-56)

James Ong

Prepared By:

BBN Systems and Technologies Corporation
10 Moulton Street
Cambridge. MA 02138

)TIC
ELECTE
MAR 20 1990
E
D
S

The views and conclusions contained in this document are those of the authors and do not represent the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Army or the United States Government.

CLEARF
FOR OPEN PUBLI TION

FEB 7 - 1990

DIRECTORATE FOR FREEDOM OF INFORMATION
AND SECURITY REVIEW (C -PA)
DEPARTMENT OF DEFENSE

90 000516

90 03 20 168

Contents

i

1. Background

Black holes can occur in packet-switched networks that use distance-vector route calculation algorithms such as tier routing. This section briefly reviews tier routing and defines new terms relevant to black holes.

1.1 Tier Routing

Packet radios use tier routing [1], a variation on distance-vector routing, to maintain routes to all radios in a non-hierarchical network or to all radios in the same cluster in a hierarchical network. Each radio's routing table contains an entry for each **destination** packet radio; the entry contains the following information:

- The destination packet radio ID

- The reporting packet radio ID

- Hop count to the destination.

    Each packet radio periodically broadcasts a **Packet Radio Organization Packet (PROP)** to all of its neighbor radios, listing its distance to every destination as stored in its routing table. Neighbor packet radios update their routing tables to incorporate shorter routes described within this PROP. For example (Figure 1.1), if radio X advertises a 3-hop route to radio Z, and radio Y is radio X's neighbor whose route to Z is 5 hops (dotted lines), radio Y will update its route when it hears X's PROP. Radio Y's new route to Z (dashed lines) will be 4 hops long and will have X as its reporting PR. Using this route, Y will forward all packets en-route to Z via neighbor radio X.


Figure 1.1: Routing Updates

1.2 Black Holes

A black hole is a packet radio whose PROPs contain false hop-count information for certain destination radios. In Figure 1.2, if radio B correctly implements the tier-routing algorithm, its PROPs will state that it is n hops from D. But, if B is a black hole, its PROPs may state that its distance to D is only k hops, where $2 \le k < n$. Although a black hole could claim that it is just one hop away from many destinations, a network management system can easily detect this "lie" by noticing that these destinations do not list the black hole as a good neighbor.

D

n-1
hops

A

B

C

E

Figure 1.2: Black Hole Example

It is likely that the false routes reported by the black hole's PROPs will be shorter than many of the correct routes known by the black hole's neighbors for the same destinations. Consequently, when neighboring radios receive the black hole's PROPs, they will update their routing tables and forward much of their traffic to the black hole. As this misinformation spreads throughout the network, much of the traffic created near the black hole will be sucked in toward the black hole, causing two problems. First, because the black hole's neighbors' routing tables contain bad routes, they repeatedly forward packets back to the black hole, and many packets never reach their destinations. Second, the excessive amount of traffic forwarded to the black hole causes it to become a traffic bottleneck. The number of lost packets

and volume of traffic drawn into the black hole will vary with the number of destinations it lied about and the amount by which the distances were understated.

Congestion around a black hole can persist despite the radios' adaptive routing algorithms. As congestion increases around the black hole, packet collisions increase, link qualities decrease, links between the black hole and its neighbors are dropped, new routes are found, and traffic is rerouted around the black hole in minutes or less. However, as the traffic around the black hole decreases, collisions will then decrease, link qualities will increase, links will be brought back up, and the black hole will again attract excessive traffic, causing congestion and delay.

Not only do packet radios report routing information to neighbors in PROPs, they report it to network management systems in MDP2 [2] packets as well. A *consistent* black hole reports the same information in its PROPs and its MDP2 packets. For example, if B is a consistent black hole, its MDP2 will report that it is k hops from D. An *inconsistent* black hole reports greater distances in its MPD2 packets than in its PROPs. For example, if B is an inconsistent black hole, its MDP2 might report that it is $k+2$ hops from radio D instead of k hops, as reported within its PROPs. An inconsistent black hole might tell a different lie to the network management system to try to avoid being detected by a consistent black hole. Although an inconsistent black hole could report lower values of distances to destinations in its MDP2 packets than it reported in its PROPs, we do not view this kind of lying to be a threat.

A black hole successfully lies to a neighbor if the neighbor updates its routing table to incorporate false short routes reported by the black hole's PROP. In our example, B successfully lied to C, and it stored $k+1$ as the distance and B as the reporting PR for its route to radio D. Since PROPs are broadcast in a single packet, a black hole must tell the same lie to all of its neighbors, but whether or not the lie successfully causes a neighbor to incorporate a false route depends upon the contents of the neighbor's routing table. In practice, we expect black holes to lie aggressively and lie successfully to most if not all of their neighbors.

A black hole lies through a neighbor by claiming that neighbor as the reporting PR for a false route in its MDP2 packet. In our example, radio B has lied through radio A if its MDP2 states that B's route to D is k hops long and that radio A is the route's reporting PR. A black hole can cause many false routes to be propagated throughout the network by lying through one, some, or all of its neighbors.

By definition, radio Y has a routing-table discrepancy with its neighbor X if:

- Y's routing table contains at least one entry with X as the reporting PR, and

- the distance value in Y's entry is less than $k+1$, where k equals X's distance to the same destination.

Routing discrepancies are not symmetric: a discrepancy between table X and table Y does not imply a discrepancy between table Y and table X.

A routing discrepancy exists when Y's distance is lower than it should be, not higher. Packet radio Y can inflict relatively little damage on the network by overstating its distances to destinations. Even if Y reports infinitely long routes to all destinations, its neighbors will simply route all their packets around it.

In the next section we outline an algorithm for detecting black holes that compares the suspected black hole's table against each of its neighbors' and compares each of the neighbors' tables against the suspected black hole's. The number of SBH $\Rightarrow$ NBR (suspected-black-hole-to-neighbor) discrepancies

equals the number of neighbors with which a suspected black hole has discrepancies. For example, Y has k SBH $\Rightarrow$ NBR discrepancies if:

- k of Y's neighbors are reporting PRs for some of Y's routes, and

- their distances for these routes are greater than Y's distances minus one.

The number of neighbor-to-suspected-black-hole (NBR $\Rightarrow$ SBH) discrepancies equals the number of neighbors that have discrepancies if k of Y's neighbors:

- have some routes with Y as the reporting PR, and

- their distances for these routes are less than Y's distances plus one.

2.   The Black-Hole-Detection Algorithm

Section 2.1 provides an overview of our black-hole-detection algorithm, which analyzes MDP2 packets from packet radios for routing discrepancies. Section 2.3 lists the type and number of discrepancies that will occur, depending upon whether or not the suspected black hole and/or its neighbors are black holes. Our algorithm correctly detects black holes if several assumptions can be made about the packet-radio network. Section 2.2 describes those assumptions. Section 2.4 describes our decision criteria that use the assumptions in Section 2.3 to classify a suspected radio based upon the routing discrepancies found between it and its neighbors.

2.1   Overview

Our algorithm accepts as input the ID of a packet radio that has been suspected of being a black hole. We imagine that a network operator or another module within the network management system might suspect the radio of being a black hole and then invoke the black-hole-detection algorithm to confirm or reject the suspicion. Our algorithm determines whether or not a suspected radio is a black hole by comparing its routing table with those of its neighbors and detecting specific patterns of inconsistency. Our algorithm has five steps:

1. Accept as input the ID of a packet radio suspected of being a black hole.

2. Query the suspect and its neighbors for MDP2 packets containing their routing information.

3. Compare the neighbors' routing tables against the suspect's tables and count the number of NBR $\Rightarrow$ SBH routing discrepancies.

4. Compare the suspect's table against those of its neighbors and count the number of SBH $\Rightarrow$ NBR discrepancies.

5. Apply the decision criteria described in Section 2.4 to classify the suspected radio as an OK radio, a consistent black hole, or an inconsistent black hole.

2.2   Routing Table Discrepancies

This section describes the number of SBH $\Rightarrow$ NBR and NBR $\Rightarrow$ SBH routing table discrepancies that will exist for a suspected radio, depending upon whether the suspect is an OK radio, a consistent-black-hole, or an inconsistent black hole.

2.2.1 Suspect is an OK Packet Radio

A suspected OK packet radio will have an NBR $\Rightarrow$ SBH discrepancy with each neighboring consistent-black-hole radio that has lied through it. For example, in Figure 1.2, let's assume that:

- B is suspected of being a black hole but is an OK packet radio, and

- C is a consistent black hole that has lied through radio B by reporting B as the reporting PR for its false route to D.

C's distance to D reported in its MDP2 will appear too low when compared to B's distance to D. Consequently, a network management system will find an NBR $\Rightarrow$ SBH discrepancy for suspected packet radio B with radio C.

A suspected OK packet radio will have an SBH $\Rightarrow$ NBR discrepancy with each inconsistent neighboring black hole that has lied to it. For example, in Figure 1.2, let's assume that:

- B is suspected of being a black hole but is an OK packet radio,

- A is an inconsistent black hole whose PROP states that A is only k hops from D where $2 \leq k < n - 1$.

- B has incorporated A's route to D, so B's MDP2 packet states that it is $k + 1$ hops from D through reporting PR A, and

- A's an inconsistent black hole, its MDP2 packet states that it's greater than k hops from D.

B's distance to D will appear too low when compared with A's distance to D. Consequently, the network management system will find an SBH $\Rightarrow$ NBR discrepancy for B with its neighbor radio A.

2.2.2 Suspect is a Consistent Black Hole

If the suspected black hole is consistent, it will contain an SBH $\Rightarrow$ NBR discrepancy with each OK neighbor that the suspect has lied through. For example, in Figure 1.2, let's assume that

- A is an OK packet radio,

- B is suspected of being a black hole and really is a consistent black hole,

- B is n hops from D, and

- B's PROP states that B is k hops from D, where $2 \leq k < n$

If B's MDP2 states that its route to destination D is k hops long through reporting PR A, B's distance to D will be too low when compared with A's distance to D.

A suspected consistent black hole will contain an SBH $\Rightarrow$ NBR discrepancy with each neighboring consistent black hole it has lied through. This case is the same as when the suspect lies through an OK packet radio.

A suspected consistent black hole will contain an NBR $\Rightarrow$ SBH routing table discrepancy for each neighboring inconsistent black hole that has lied to the suspect. For example, in Figure 1.2, let's assume that

- A is an inconsistent-black-hole neighbor radio that has reported to suspected black-hole radio B that A is m hops from destination D, $m < n - 1$.

- B is a suspected as a black hole and really is a consistent black hole.

- Consistent-black-hole radio B has incorporated the false route advertised by A's PROP and reports in its own PROP and MDP2 packet that it is $< m + 1$ hops away from D, and

- Radio A has reported in its MDP2 packet that A is $> m$ hops away from D.

Consequently, B's distance will appear too low when compared with A's distance to D reported in its MDP2 packet.

2.2.3 Suspect is an Inconsistent Black Hole

An inconsistent black hole will have an NBR $\Rightarrow$ SBH discrepancy with each OK packet radio it has lied to. For example, in Figure 1.2, let's assume that:

- B is suspected as a black hole and really is an inconsistent black hole.

- B reports in its PROP that it is k hops from D, where $2 \leq k < n$.

- B reports in its MDP2 that is is greater than or equal to $k + 1$ hops from D.

- E is an OK packet radio, and

- B has successfully lied to all of its neighbors.

Packet radio E's MDP2 packet will report a route to D that is k+1 hops long through reporting PR B. As an inconsistent black hole, B's MDP2 packet will claim that B is greater than or equal to $k + 1$ hops from D. Consequently, E's distance to D will appear too low when compared to reporting PR B's distance to D.

Radio B will also have an NBR $\Rightarrow$ SBH discrepancy with each consistent black hole it has lied to. This is basically the same case as a neighboring OK packet radio.

The routing discrepancies that may exist between two inconsistent black holes depend upon the details of how the radios have lied in their PROP and MDP2 packets.

Table 2.1 summarizes the discrepancies detected in each of these situations.

2.3 Assumptions

Our algorithm described in Section 2.1 correctly classifies suspected radios into OK radios, consistent black holes, and inconsistent black holes if certain assumptions can be made about the packet radio network.

Suspected Radio	MDP2 Routing Discrepancies
OK Packet Radio	NBR ⇒ SBH discrepancy with each consistent black hole that lied through the suspect
	SBH ⇒ NBR discrepancy with each inconsistent black hole that lied to the suspect
Consistent Black Hole	SBH ⇒ NBR discrepancy with each OK packet radio the suspect lied through
	SBH ⇒ NBR discrepancy with each consistent black hole the suspect lied through
	NBR ⇒ SBH discrepancy with each inconsistent black hole that lied to the suspect
Inconsistent Black Hole	NBR ⇒ SBH discrepancy with each OK radio the suspect lied to
	NBR ⇒ SBH discrepancy with each consistent black hole the suspect lied to
	NBR ⇒ SBH or SBH ⇒ NBR discrepancy with each inconsistent black hole

Table 2.1: Routing-Table Discrepancies

1. **All routing inconsistencies are caused by black holes**

   Our algorithm assumes that all routing inconsistencies are caused by lies told by black-hole packet radios. Routing inconsistencies caused by other factors such as changing network topology can confuse the algorithm. Dependence on this assumption seriously limits the applicability of our algorithm.

2. **There is only one black hole in any local area**

   It is harder to tell which radios are lying and which are telling the truth as the density of liars increases in a region of the network. Our algorithm assumes that:

   - Each OK radio has at most one black hole as neighbor, and
   - Each black hole radio has only OK radios as neighbors.

   Dependence on this assumption somewhat limits the applicability of our algorithm.

3. **Each black hole successfully lies to at least two neighbors**

8

Our algorithm assumes that if a black hole understates its distances to destinations, at least two of its neighbors will incorporate these false routes into their tables. This assumption is plausible.

2.4 Decision Criteria

This section describes how the algorithm classifies a suspected packet radio, given the number of discrepancies found between the suspect and its neighbors. If the assumptions described in Section 2.3 are true, this decision criteria will correctly classify a radio as one of:

- OK packet radio.

- Consistent black hole, or

- Inconsistent black hole.

2.4.1 Two or More NBR ⇒ SBH Discrepancies

An NBR ⇒ SBH discrepancy indicates that either:

- the suspect is an inconsistent black hole that has lied to an OK neighboring packet radio, or

- the suspect is an OK packet radio, but the neighbor is a consistent black hole that has lied through the suspect.

Since we have assumed that each packet radio has at most one neighboring black hole, finding two or more NBR ⇒ SBH discrepancies indicates that the suspect is an inconsistent black hole.

2.4.2 Exactly One NBR ⇒ SBH Discrepancy

We have assumed that every black hole successfully lies to at least several of its neighbors, so we should detect several NBR ⇒ SBH discrepancies, not just one, if the suspect is an inconsistent black hole. Consequently, detection of just a single NBR ⇒ SBH discrepancy leads us to believe that the suspect is an OK packet radio and that the neighbor is a consistent black hole.

2.4.3 Two or More SBH ⇒ NBR Discrepancies

An SBH ⇒ NBR discrepancy indicates that either:

- the suspect is a consistent black hole that has lied through the neighboring OK packet radio, or

- the suspect is an OK packet radio, but the neighbor is an inconsistent black hole that has lied to the suspect.

Since we have assumed that each packet radio has at most one neighboring black hole, finding two or more SBH ⇒ NBR discrepancies indicates that the suspect is a consistent black hole.

2.4.4   Exactly One SBH ⇒ NBR Discrepancy

If just a single SBH ⇒ NBR discrepancy exists, we need additional information to disambiguate among the possibilities:

- The original suspect is a consistent black hole and the neighbor is an OK packet radio, or

- the neighbor is an inconsistent packet radio and the original suspect is an OK packet radio.

Let's call the neighboring radio N. If we compare N's routing information to those of its neighbors and count the number of SBH ⇒ NBR and NBR ⇒ SBH discrepancies between N and its neighbors, we can obtain the necessary additional information. If N is an OK packet radio (first case), we should see just a single discrepancy between N and the original suspect. If N is an inconsistent black hole (second case), we should see several NBR ⇒ SBH discrepancies between N and its neighbors, one for each neighbor that N lied to.

Table 2.2 summarizes the decision criteria that classify packet radios based upon the number and type of routing discrepancies found.

Discrepancies	Classification
Two or More NBR ⇒ SBH	Suspect is an inconsistent black hole that has lied to these neighbors.
Exactly One NBR ⇒ SBH	Suspect is an OK radio. Neighbor is a consistent black hole that has lied through the suspect.
Two or More SBH ⇒ NBR	Suspect is a consistent black hole that has lied through these neighbors.
Exactly One SBH ⇒ NBR	If the neighbor has more than one NBR ⇒ SBH discrepancy with *its* neighbors, the original suspect is an OK radio and the neighbor is an inconsistent black hole. Otherwise, the suspect is a consistent black hole that has lied through this one neighbor.

Table 2.2: Black Hole Decision Criteria

3. Algorithm Implementation

3.1 Simulation

We built a simulation of a packet-radio network containing black holes. The simulation creates networks of packet radios at random locations and considers two radios to be neighbors if the distance between them falls below a user-specified threshold. The simulated radios exchange PROPs and update their tables to incorporate shorter routes reported by their neighbors[1]. The user may specify any of the radios to be consistent or inconsistent black holes that create and send PROPs and MDP2s containing false information. This simulation is build on top of a tool written in Common Lisp and Flavors that enables simulation of packet radio networks to be written more easily [3].

3.2 Interactive Graphics Display

We built an interactive graphics display that displays a circular icon for each packet radio and a line-shaped icon for each good link between two packet radios. In addition to showing the network's connectivity, a person can click the mouse over a packet-radio icon and receive either:

- a textual display of information about a packet radio such as its routing table, PROP packet, MDP2 packet, routing discrepancies with its neighbors.

- a color-coding of the packet radios' icons in the display based upon their distances from a specified radio according to the specified radio's routing table, PROP, or MDP2.

This interactive display helped provide insights into packet-radio routing, especially in the presence of black holes.

3.3 Classification Algorithm

Using the interactive display, the user can invoke the black-hole-detection algorithm on any simulated radio by mouse-clicking on its icon. Our implementation of the algorithm then retrieves MDP2 information from the specified simulated radio and its neighbors, and classifies the radio.

---

[1]The simulation currently does not simulate the loss-of-neighbor connectivity or the dropping of routes.

4. Summary

We have designed and implemented an algorithm that detects black holes in simulated packet-radio networks. The algorithm works even if the black hole lies inconsistently, but it relies on several optimistic assumptions. For example, it incorporates the effects of routing discrepancies normally caused by changing network topology. Also, the algorithm assumes that a network manager can request MDP2 packets from radios near a black hole, even though it may be difficult for packets to leave such an area. Dependence on these simplifications prevents the algorithm from detecting black holes in real packet-radio networks.

Modification of the packet-radio algorithms can address the problem of black holes more directly. For example, networks which use shortest path first (SPF) routing rather than distance vector routing avoid black holes entirely. Even if distance vector routing is used, we could modify PROPs to include the reporting PR for each routing entry. Neighbor radios of the PROP's source could detect faulty routes in the PROP that lie through them. However, if a radio claims that its neighbor is a black hole, the network manager cannot know which radio to believe without additional information. Also, this last technique does not detect a black hole that advertises within its PROP a non-existent or distant radio as a reporting PR.

Our main objective in studying black holes was to explore and model the routing characteristics of packet radios by studying a single, well-defined problem. Our work in black holes forms a basis for future modeling of packet-radio routing behavior.

Bibliography

[1] J. Westcott and J. Jubin. A distributed routing design for a broadcast environment. In *Milcom '82*. pages 10.4.1–10.4.5. Boston. MA, 1982.

[2] M. Cote. *Monitoring CAP8 and SURAP Networks*. SURAN Program Technical Note Number 12. BBN-STC. Boston, MA, 1989.

[3] J. Ong and M. Cote. Simulation and scheduling tools for SURAN. 1987. Send requests to Gregory Lauer. BBN-STC, 10 Moulton Street. Cambridge, MA 02138.